Virtual Switches Demand Rethinking Connectivity for Servers

With the creation of multiple virtual machines on each physical host server, network adapters now connect servers to the network with equivalency to a traditional top-of-rack switch.

The adoption of virtualization has led to a paradigm shift in the data center. At the core of these changes is the fact that a single host server running multiple virtual machines (VMs) can now represent the equivalent of multiple physical servers. Traffic among the VMs and with the outside world is now managed and coordinated by a virtual switch.

The virtual switch exists in software built into the hypervisor, meaning that it resides within the server itself, rather than as a separate physical device. Therefore, the network connectivity that was situated between the servers and a top-of-rack switch now occurs within the host server. By extension, the host server's network adapters now handle the connections that formerly existed between the top-of-rack switches and core switches.

This shift can be further described as follows:

- **Yesterday's data center: Physical edge switches.** In conventional, non-virtualized data centers, servers were connected to top-of-rack switches by multiple Gigabit Ethernet (GbE) connections, and those switches connected to core switches with large trunked uplinks.
- **Virtualized data centers: Virtual edge switches within servers.** Today's virtualized host servers process network traffic in software, with the hypervisor's virtual switch replacing the top-of-rack switch for server connectivity, and the trunked uplinks are handled by the host server's Ethernet adapters.
- **Supporting the new paradigm: Implications for networking.** Each physical host server in a virtualized environment is the logical equivalent of a rack in the non-virtualized one, meaning that each physical host server's network connection must actually be treated like a trunked uplink found on switches.

In the typical virtualized environment, data is segregated within trunked connections using multiple virtual LANs (VLANs). This approach provides logical separation between the data streams that enables them to be prioritized and secured as needed. VLANs have become a preferred topology in many mainstream implementations.

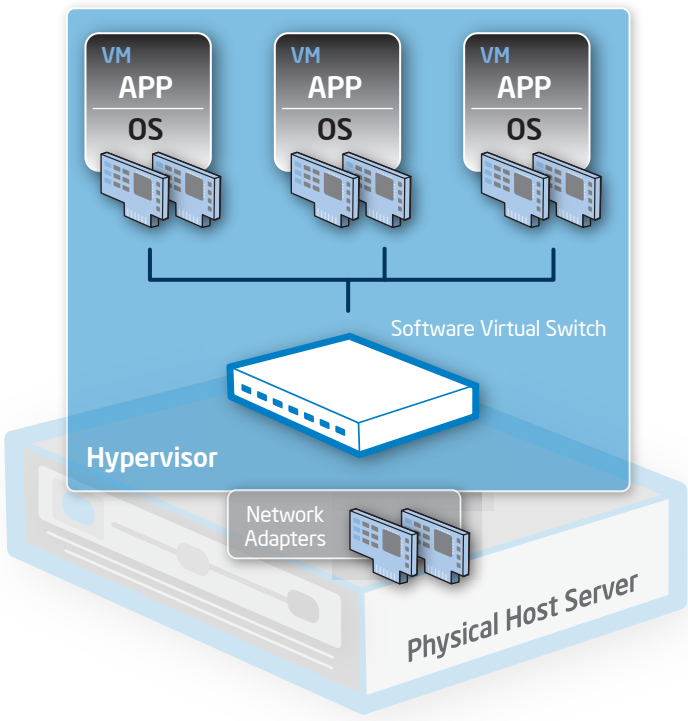

Figure 1. In a virtualized environment, the hypervisor running on the physical host server includes a software-based virtual switch.

Yesterday's Data Center: Physical Edge Switches

In a traditional data center where virtualization is not used, each server is connected to a top-of-rack switch, typically with two GbE connections. These network connections can provide redundancy for fault tolerance, or they can be aggregated to provide a total of 2 Gb bandwidth to the server. In many cases, to provide greater redundancy, the two connections lead to separate switches.

The top-of-rack switches, in turn, are uplinked to one or more core switches:

- **Multiple servers per connection.** Trunk connections from the top of rack switches each carry traffic between multiple servers in the rack and the core switch infrastructure.
- **High bandwidth demands.** In a typical implementation, switch-to-switch connections are based on 10 Gigabit Ethernet (10GbE).

This familiar topology is an effective means of centrally connecting a potentially large number of servers to a core switch infrastructure. As the environment grows, however, the large number of server adapters and switches can carry significant capital and operating expense.

Virtualized Data Centers: Virtual Edge Switches within Servers

Using server virtualization, the number of physical host servers required can be reduced. Multiple VMs, each representing what used to be a distinct physical server, can be placed on a single host server and started and stopped independently, and they can even use different operating systems. They can be provisioned and brought online quickly and can be load balanced in an automated fashion if desired. They also provide a high degree of fault tolerance. As a result, IT can realize significant cost and TCO savings by embracing virtualization.

Since the VMs running on a single physical host server need to talk to each other or other systems outside of the host server, the hypervisor must also provide network switch functionality. That switching is analogous to the top-of-rack switches in the non-virtualized data center.

- **Each host server contains multiple virtualized servers.** Now a single physical host server represents a traditional rack of servers.

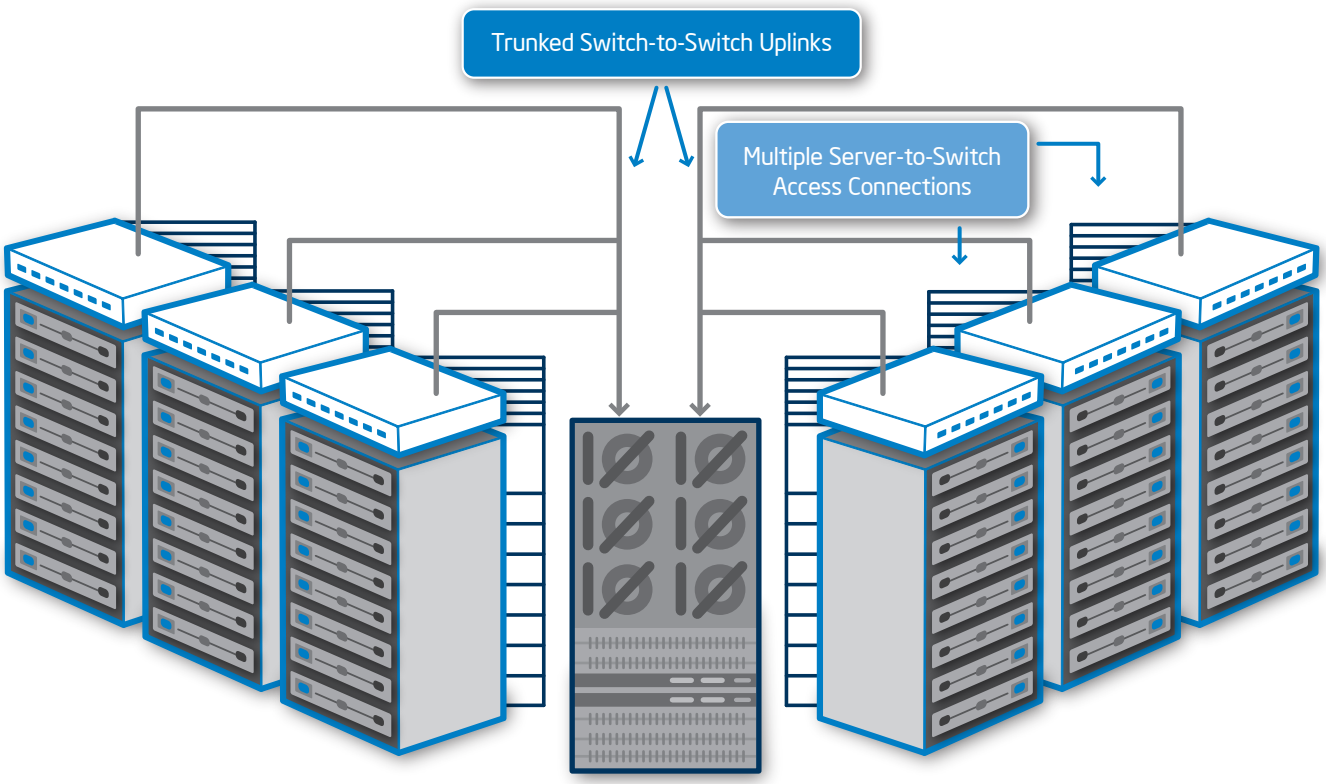

Figure 2. Non-virtualized servers are connected to top-of-rack switches, which are connected to core switches using large trunked uplinks.

- **Server adapters handle trunked traffic.** Each physical network connection is now an uplink from a host server's virtual switch to a physical switch, handling traffic from multiple VMs to the outside systems. Network and Ethernet-based storage traffic are separated and isolated from each other through the use of VLANs. The uplinks then trunk all the different VLANs on a single physical connection, which allows all authorized traffic to move between the virtual switch and the physical switch.

The reduced number of physical connections (uplinks) under this virtualized model requires fewer server adapters and switch ports than in the non-virtualized case, helping to save on both capital and operating expenses. The high bandwidth requirements for each connection benefit dramatically from using 10GbE server adapters, rather than using multiple GbE connections.

Supporting the New Paradigm: Implications for Networking

Today's mainstream data center servers based on multicore Intel® Xeon® processors provide very high levels of computing capacity. That performance headroom enables increasingly large numbers of VMs per physical host. As that number grows, providing two GbE server adapters for each VM quickly becomes untenable.

Upgrade to 10 Gigabit Ethernet

Rather than using unwieldy numbers of GbE connections, the virtualized model clearly benefits from a smaller number of 10GbE connections. Consolidating onto 10GbE trunks can also provide better bandwidth utilization and help reduce cabling requirements and electrical costs. Networking in the virtualized environment follows logically from that in the non-virtualized case, while also providing additional functionality:

- **VLAN traffic segmentation.** Data streams can be differentiated logically using the switch's VLAN functionality. As a very secure, proven technology, VLANs provide isolation capabilities that are comparable to those of separate physical links.
- **Quality of service.** The virtual switches built into leading hypervisors provide increasingly sophisticated quality-of-service functionality, helping to ensure that specific amounts of bandwidth are available for various types of traffic, effectively prioritizing it as needed.

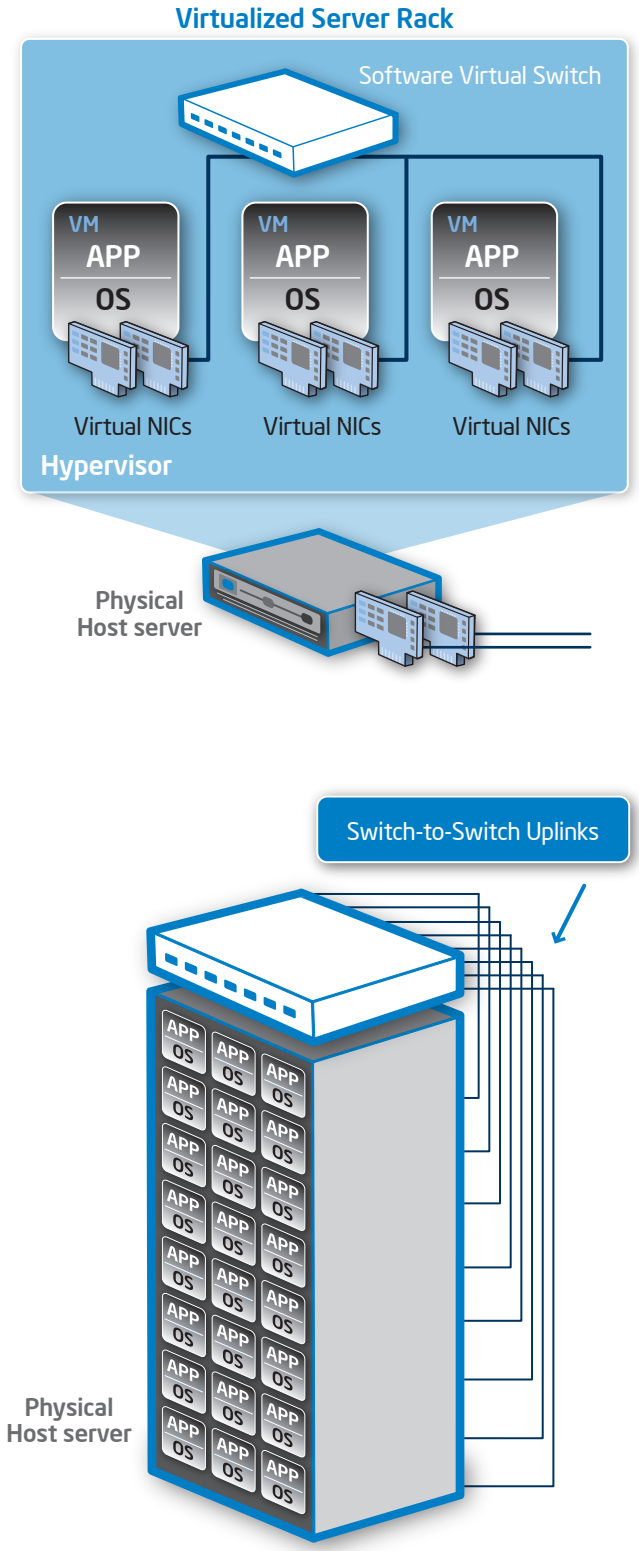

Figure 3. Virtualized servers are connected to switches with trunked uplinks equivalent to non-virtualized switch-to-switch connections.

Take Advantage of Virtual Switching Technologies from Intel

Various technologies that Intel has developed (or helped develop) help reduce the software overhead associated with virtual switching. Some of the more important of these include the following:

- Virtual Machine Device Queues (VMDq)[1] offloads part of the burden of processing network traffic from the hypervisor to the network controller, enabling 10Gb connections to deliver near-native throughput.
- Single Root I/O Virtualization (SR-IOV)[2] is an industry specification supported by Intel that allows multiple VMs to access network hardware directly, helping dramatically improve virtualized performance and direct assignment scalability.
- Intel® Virtualization Technology for Directed I/O (Intel® VT-d)[3] allows a VM to have direct, exclusive access to an I/O device and system memory, helping to provide increased memory protection through isolation.

Getting the full benefit of virtualization on today's powerful multicore Intel® processor-based servers requires the bandwidth advantages of 10GbE server adapters working in conjunction with Intel® Virtualization Technology for Connectivity (Intel® VT-c). By delivering advanced technologies for virtualization, Intel® Ethernet 10 Gigabit server adapters help forward-looking IT organizations continue expanding services while contributing to a healthier bottom line.

Learn More About Intel® Ethernet: **www.intel.com/go/ethernet**

[1] http://www.intel.com/network/connectivity/vtc_vmdq.htm
[2] http://download.intel.com/design/network/applnots/321211.pdf
[3] http://download.intel.com/technology/computing/vptech/Intel(r)_VT_for_Direct_IO.pdf

 0910/BY/MESH/PDF 324253-001US